

REMARKS

Claims 1, 6, 7, 13, 14, 17, 18, 19, 21, 26, 27, 32, 33, and 37-42 have been amended. Claims 3-5, 16, 23-25, and 34-36 have been canceled. Claims 1-2, 6-15, 17-22, 26-33, and 37-42 are pending in the application. Reconsideration is respectfully requested in light of the following remarks.

Section 112, Second Paragraph, Rejections:

The Examiner rejected claims 5, 25, and 36 under 35 U.S.C. § 112, second paragraph, for insufficient antecedent basis for the limitation “the storage class”. Applicants have canceled claims 5, 25, and 36. Therefore, removal of the § 112 rejection is respectfully requested.

The Examiner rejected claims 7, 18, 27, and 38 under 35 U.S.C. § 112, second paragraph, for insufficient antecedent basis for the limitation “the write-locked storage class”. Applicants have amended claims 7, 18, 27, and 38 to recite “a write-locked storage class.” Therefore, removal of the § 112 rejection of claims 7, 18, 27, and 38 is respectfully requested.

Section 102(b) Rejection:

The Examiner rejected claims 1-5, 8-16, 19-25, 28-36 and 39-42 under 35 U.S.C. § 102(b) as being anticipated by Kishi (U.S. Patent 6,029,179). Applicants respectfully traverse this rejection for at least the reasons below.

In regard to claim 1, Applicants respectfully assert that the cited art does not teach the invention as disclosed in claim 1 of the present application as originally provided. However, Applicants have amended claim 1 for further clarity.

In regard to claim 1, the cited art does not teach or suggest *file system software configured to assign and migrate data in a multi-class file system comprising a hierarchy of storage classes, wherein migrated data remains online within the multi-class file system.*

Applicants note that Kishi teaches a “method and system are described for recovering data from an unreliable tape within a **tape server** which includes a **plurality of tape drives**, a **direct access storage device**, and a storage manager” (Abstract). In the Background section (column 1, lines 11-45), parts of which were cited by the Examiner in the rejection of claim 1, and relevant parts of which are cited below, Kishi describes hierarchical storage systems, and further describes, as an example of a hierarchical storage system, a virtual tape storage (VTS) system. The underlined parts are further discussed below to highlight the distinction between what Kishi teaches and what is disclosed in claim 1 of the present application:

In hierarchical storage systems, intensively used and frequently accessed data is stored in fast but expensive memory. One example of a fast memory is a direct access storage device (DASD). In contrast, less frequently accessed data is stored in less expensive but slower memory. Examples of slower memory are tape drives and disk drive arrays. The goal of the hierarchy is to obtain moderately priced, high-capacity storage while maintaining high-speed access to the stored information.

One such hierarchical storage system is a virtual tape storage system (VTS), including a host data interface, a DASD, and a number of tape devices. When the host writes a logical volume, or a file, to the VTS, the data is stored as a file on the DASD. Although the DASD provides quick access to this data, it will eventually reach full capacity and a backup or secondary storage system will be needed. An IBM 3590 tape cartridge is one example of a tape device that could be used as a backup or secondary storage system.

When the DASD fills to a predetermined threshold, the logical volume data for a selected logical volume, typically the oldest, is **removed from the DASD** to free space for more logical volumes. The selected DASD file is then appended onto a tape cartridge, or a physical volume, with the original left on the DASD for possible cache hits. When a DASD file has been appended to a tape cartridge and the original remains on the DASD, the file is "premigrated".

When the host reads a logical volume from the VTS, a cache hit occurs if the logical volume currently resides on the DASD. If the logical volume is not on the DASD, the storage manager determines which of the physical tape volumes contains the logical volume. The corresponding physical volume is then mounted on one of the tape devices, and the data for the logical volume is transferred back to the DASD from the tape.

Applicants further note that the Background section of the present application also describes the operation of such Hierarchical Storage Management (HSM) systems. Again, the underlined parts are further discussed below to highlight the above-cited distinction between what Kishi teaches and what is disclosed in claim 1 of the present application:

Rather than making copies of files as in a backup system, HSM migrates files to other forms of storage, freeing hard disk space. Events such as crossing a storage threshold and/or reaching a certain file “age” may activate the migration process. As files are migrated off primary storage, HSM leaves stubs to the files on the hard drive(s). These stubs point to the location of the file on the alternative storage, and are used in automatic file retrieval and user access. The stub remains within the file system of the primary storage, but the file itself is migrated “offline” out of the file system onto the alternative storage (e.g. tape).

In HSM, when a file that has been migrated to a lower rank of storage, such as tape, is accessed by an application, the stub may be used to retrieve and restore the file from the lower rank of storage. The file appears to be accessed by the application from its initial storage location, and demigration of the file back into the file system is performed automatically by the HSM system using the stub. While on the surface this demigration may appear transparent to the user, in practice the process of accessing and restoring the file from offline storage (e.g. tape) may introduce a noticeable time delay (seconds, minutes, or even hours) to the user when compared to accessing files stored on primary storage. Thus, accessing offloaded data in an HSM system is typically non-transparent to the application or user because of the difference in access time. In addition, since HSM introduces a substantial time lag to access offloaded data, HSM systems typically only offload low access (essentially, no access) data.

As noted, the above citations from the cited reference and the present application highlight a distinction between what Kishi teaches and what is recited in claim 1 of the present application. **Namely, this distinction is that a hierarchical storage system,**

such as the VTS (also referred to by Kishi as the “tape server”) disclosed by Kishi, when migrating files from a DASD to “a tape cartridge, or a physical volume”, migrates files/data offline (i.e., the data is no longer “online” in the file system), whereas claim 1 of the present application discloses a multi-class file system wherein data migrated from a first storage class to a second storage class remain online within the file system. In the “tape server” disclosed by Kishi, files migrated to tape cartridges are migrated offline and thus out of a file system. As disclosed by Kishi, all that remains online within the file system is information via which the storage manager may determine which of the physical tape volumes contains a requested logical volume (that is stored to tape, and thus is not online). Kishi discloses that a logical volume is migrated offline to a tape cartridge, while only location information for the logical volume (i.e., which tape cartridge the logical volume is stored on) that the storage manager may use to retrieve the logical volume, if needed, remains online in the file system.

The last paragraph in the citation from Kishi’s background section makes it clear that Kishi is directed at a “tape server” that migrates data offline:

When the host reads a logical volume from the VTS, a cache hit occurs if the logical volume currently resides on the DASD. If the logical volume is not on the DASD, the storage manager determines which of the physical tape volumes contains the logical volume. The corresponding physical volume is then mounted on one of the tape devices, and the data for the logical volume is transferred back to the DASD from the tape.

Furthermore, in the “tape server” as disclosed by Kishi that migrates files/data offline to tape cartridges, the files/data (“logical volumes”) are not directly readable or modifiable by “users” while on the tape cartridges (and thus offline). As Kishi discloses in the above citation, any access, even a read access, initiated by a user on a logical volume that has been migrated to a tape cartridge from the DASD requires that the logical volume be retrieved from a tape cartridge and moved back to the DASD. The access request is then fulfilled from the short-term storage media. Further, retrieving a migrated, offline logical volume for a user may require physical media (a tape cartridge) to be loaded into a tape device: “The corresponding physical volume is

then mounted on one of the tape devices, and the data for the logical volume is transferred back to the DASD from the tape.

Applicants note that the present application, as made clear in the currently modified claim 1, discloses that data migrated in the multi-class file system remains online in the multi-class file system. The migrated data, while not modifiable by applications while on one or more of the storage classes, is readable by applications (users) while on these storage classes. Thus, data migrated to any storage class in the multi-class file system **remains online** within the multi-class file system and is thus accessible to applications (users) via the file system software for read operations, but not for write operations. The data is not migrated offline, and thus does not need to be moved from one storage class to another storage class for read-only access by the applications. Nor does migrated data need to be “transferred back to” a storage class for read/write access. Nor does migrating data or accessing migrated data in the multi-class file system require a physical volume to be **mounted on a tape device**.

In further regard to claim 1, the cited art does not teach or suggest an application configured to *perform an operation on data stored in the multi-class file system, wherein the operation requires stable data, wherein, to perform the operation, the application is configured to: perform the operation on at least one of the one or more storage classes that store data that is not modifiable without using a split mirror of the at least one of the one or more storage classes; and perform the operation on at least one of the one or more other storage classes that store data that is modifiable using a split mirror of the at least one of the one or more other storage classes.*

In regard to the limitation *performing the operation on at least one of the one or more storage classes that store data that is not modifiable without using a split mirror*, the Examiner asserts “Kishi discloses this limitation as ‘The first step in read-only processing performed by the automated administrator 36 is to query the storage manager 30 for a list of all the logical volumes or files that the storage manager 30 has written to the physical volume at step 42.’ (Col 5, lines 19-22).” This citation simply discloses that

the automated administrator may query the storage manager for a list of all logical volumes or files that the storage manager has written to a particular physical volume. The physical volume itself is not directly involved; the information, as disclosed by Kishi, is maintained elsewhere by the storage manager. In other words, this citation does not describe performing an operation on the physical volume.

In regard to the limitation *performing the operation on at least one of the one or more other storage classes that store data that is modifiable using a split mirror*, in regard to claim 3, now canceled, which recited a limitation similar to the above, the Examiner asserts that Kishi discloses this limitation as “Another data protection scheme maintains a redundant set of the logical volumes that are stored on each tape. Each time a volume is written from the DASD to a tape device, a second copy of the volume could be mirrored on another tape. Then, if one tape device becomes unreliable, the storage manager may determine a list of volumes that were on the unreliable tape device, and then determine the location of the mirror copy corresponding to each of the volumes on the unreliable tape.” (Kishi, col. 2, lines 20-29). Applicants first note that the citation is from the Background section of Kishi, and is describing a conventional “data protection scheme” which Kishi’s disclosed “tape server” is intended to overcome the limitations of. The next paragraph in Kishi’s Background section states:

Full copying of an unreliable tape, backup methods, and **redundant storage methods** all employ considerable space on tape devices to carry out. Methods for protecting volume on tape devices are needed that do not use excessive amounts of storage space or require additional tape devices on the VTS.

Kishi portrays the “data protection scheme” described in col. 2, lines 20-29 as a “redundant storage method”. Such a “redundant storage method” is not intended for anything like *performing an operation on at least one of the one or more other storage classes that store data that is modifiable using a split mirror*. As the citation indicates, in the redundant storage method, copies of “volumes” are created and maintained as part of a data protection scheme. The citation is merely describing keeping backups or redundant copies of data (“volumes”) so that, if one copy is damaged, another copy may be used to recover the data. The cited selection, nor anywhere else in Kishi, does not

teach or suggest such a limitation.

Furthermore, as noted above, Kishi's disclosed "tape server" is intended to overcome the limitations of the cited "data protection scheme". Kishi states that such redundant storage methods employ "considerable space on tape devices to carry out" and that "methods for protecting volume on tape devices are needed that do not use excessive amounts of storage space or require additional tape devices on the VTS." Thus, Kishi, in disclosing Kishi's "tape server", actually appears to **teach away** from keeping "mirror copy corresponding to each of the volumes on the unreliable tape".

In any case, the notion of a "mirror copy corresponding to each of the volumes on the unreliable tape", as disclosed in Kishi's background section in regard to a conventional "data protection scheme" described as a "redundant storage method", is nothing like the notion of using a split mirror in performing an operation on a storage class that stores data that is online and that is modifiable, as is disclosed in claim 1.

Contrary to the Examiner's assertions, Kishi clearly does **not** teach or suggest the above limitations as recited in claim 1 in the cited selections from Kishi. However, even if Kishi **did** teach the limitations in the citations provided by the Examiner, nowhere does Kishi teach or suggest what is recited in claim 1 **when the claim is considered as a whole**. Kishi does not teach or suggest an *application configured to perform an operation on data stored in the multi-class file system, wherein the operation requires stable data, wherein, to perform the operation, the application is configured to: perform the operation on at least one of the one or more storage classes...without using a split mirror; and perform the operation on at least one of the one or more other storage classes...using a split mirror.*

Applicants remind the Examiner that anticipation requires the presence in a single prior art reference disclosure of each and every element of the claimed invention, arranged as in the claim. M.P.E.P 2131; *Lindemann Maschinenfabrik GmbH v. American Hoist & Derrick Co.*, 221 USPQ 481, 485 (Fed. Cir. 1984). The **identical invention**

must be shown in as complete detail as is contained in the claims. *Richardson v. Suzuki Motor Co.*, 9 USPQ2d 1913, 1920 (Fed. Cir. 1989). **As shown above, Applicants' claimed invention is clearly not anticipated by Kishi.**

Thus, for at least the reasons presented above, the rejection of claim 1 is not supported by the cited art and removal thereof is respectfully requested. Similar remarks as those above regarding claim 1 also apply to claims 14, 19, 21, and 32.

In regard to claim 13, the arguments given above in regard to the rejection of claim 1 apply equally to claim 13. Thus, for at least the reasons presented above, the rejection of claim 13 is not supported by the cited art and removal thereof is respectfully requested.

Section 103(a) Rejection:

The Examiner rejected claims 6, 7, 17, 18, 26, 27, 37 and 38 under 35 U.S.C. § 103(a) as being unpatentable over Kishi in view of Thomas et al. (U.S. Patent 6,061,692) (hereinafter "Thomas"). Since the rejections have been shown to be unsupported for the independent claims, a further discussion of the § 103(a) rejections is not necessary at this time.

In regard to the rejections under both § 102 and § 103, Applicants also assert that numerous ones of the dependent claims recite further distinctions over the cited art. However, since the rejections have been shown to be unsupported for the independent claims, a further discussion of the dependent claims is not necessary at this time.

CONCLUSION

Applicants respectfully submit that the application is in condition for allowance, and prompt notice to that effect is respectfully requested.

If any fees are due, the Commissioner is authorized to charge said fees to Meyertons, Hood, Kivlin, Kowert, & Goetzel, P.C. Deposit Account No. 501505/5760-15100/RCK.

Respectfully submitted,

/Robert C. Kowert/

Robert C. Kowert, Reg. #39,255
Attorney for Applicants

Meyertons, Hood, Kivlin, Kowert, & Goetzel, P.C.
P.O. Box 398
Austin, TX 78767-0398
Phone: (512) 853-8850

Date: July 10, 2007